Online version at https://journal.lenterailmu.com/index.php/jafotik

PT. Lentera Ilmu Publisher

JAFOTIK

Jurnal Sistem Informasi dan Teknik Informatika

E-ISSN: 3031-2698, DOI prefix: 10.70356

Design and Implemen tation of a Computer Network to Improve Data Center Accessibility Using Arista Devices and a Failover Method at Terminal 3 of Soekarno-Hatta Airport

Elang Nurjaya [1], Akrom [1]*

[1] *Informatics Engineering Study Program, Faculty of Computer Science, Universitas Pamulang, Tangerang Selatan, Indonesia*

ARTICLE INFO

Article history:
Received 01 August 2026
Revised 02 September 2026
Accepted 12 September 2026
Published 15 September 2026

Keywords:
Data Center
Spine-Leaf
VXLAN
BGP EVPN
MLAG
Failover

ABSTRACT

Data center network accessibility and availability are critical to supporting reliable data services, particularly in environments requiring continuous connectivity. This study aims to design and implement a redundant computer network at Terminal 3 of Soekarno-Hatta International Airport using Arista devices and a failover mechanism. The proposed architecture adopts a Spine-Leaf topology with VXLAN as the data plane, BGP EVPN as the control plane, and MLAG, Anycast VTEP, and ECMP to provide redundancy and multiple active forwarding paths. The implementation was conducted using Arista vEOS 4.31.0F in a PNETLab simulation environment. Network performance was evaluated through three failure scenarios: access-link failure, Leaf device failure, and Spine device failure. The results showed 0% packet loss in all scenarios, with network recovery measured at less than 200 ms. The Leaf failure caused an approximately 20-second HTTP service disruption, while the Spine failure caused a 130-second disruption. Further analysis indicated that the prolonged Spine failure disruption was caused by different BGP timer configurations between the Arista fabric and FortiGate perimeter, rather than by fabric convergence, which occurred within 9 seconds. The proposed architecture successfully improved network redundancy, resilience, and data center accessibility.

1. Introduction

Terminal 3 of Soekarno-Hatta International Airport is an operational environment that requires computer network infrastructure with a high level of availability and reliability. Various operational services, such as the Airport Operational Database (AODB), security systems, automated baggage systems, Closed-Circuit Television (CCTV), e-gates, and internal communications, depend on a stable, secure, and low-latency network. The design of the Data Center network architecture is

* *Corresponding author:* Akrom
E-mail address: dosen02613@unpam.ac.id

therefore an important factor in maintaining service continuity when failures occur in network devices or communication links [1]. Terminal 3 of Soekarno-Hatta International Airport has high data traffic requirements and supports various operational services. The network infrastructure must provide stable connectivity to support data exchange among systems and devices. As service requirements and network traffic continue to increase, the infrastructure needs to provide redundancy, optimal utilization of available paths, and an effective recovery mechanism when network failures occur [2].

The main problem identified in the network infrastructure used as the research object is the use of a single Backbone Link without a backup path. This condition creates a high dependency on a single communication path. In addition, the implementation of the Spanning Tree Protocol (STP) causes several Uplink paths to operate in a blocking state. Consequently, not all available network capacity can be actively utilized. This condition also increases the risk of connectivity disruption when the primary communication path experiences a failure [3].

These conditions indicate the need for a network design that can provide redundant paths while simultaneously utilizing multiple communication paths actively. Redundancy is required to ensure that a failure in one path or device does not immediately interrupt network connectivity. Furthermore, the active utilization of multiple paths is required to optimize the use of available network capacity [4].

The Leaf-Spine architecture is used as the main approach for the network design in this research. This architecture consists of Spine and Leaf layers that are interconnected to provide multiple communication paths. Each Leaf device is connected to multiple Spine devices, allowing alternative paths to remain available when one of the paths fails. This structure also enables the implementation of Equal-Cost Multi-Path (ECMP), allowing traffic to be distributed across multiple paths with equal routing costs.

In this research, two Spine devices and four-Leaf devices are used. The Leaf devices serve as part of the service and access layers, while the Spine devices form the fabric layer that provides communication paths between network layers. This design changes the previous network condition, which relied on a single Backbone path, into a network architecture with communication paths through two Spine devices. Therefore, if one Spine device fails, an alternative communication path remains available through the other active Spine device [5].

To improve redundancy at the access layer, this research implements Multi-Chassis Link Aggregation (MLAG). MLAG allows two Leaf devices to operate as a logical pair and provide Active-Active connectivity to the Switch Access devices. Through this mechanism, the access devices are not dependent on a single Leaf device or a single Uplink path. In addition to providing device and link redundancy, the configuration allows both Uplink paths to be actively utilized, resulting in more efficient use of network capacity [6].

In addition to MLAG, an Anycast Gateway is implemented to provide the same gateway function across the Leaf pair. This configuration supports traffic redirection when one of the Leaf devices experiences a failure. By providing redundant gateway functionality across the Leaf pair, client connectivity can be maintained when one of the Leaf devices becomes unavailable. At the network fabric layer, ECMP is implemented to provide multiple paths toward destination devices. This mechanism allows traffic to utilize more than one path with equal routing costs. The implementation of ECMP at the Spine layer is one of the main mechanisms evaluated when a Spine device fails. With two Spine devices available, traffic can be redirected through the remaining active Spine device when one of the Spine devices experiences a failure.

To support network segmentation within the Data Center environment, this research implements Virtual Extensible LAN (VXLAN) and Border Gateway Protocol Ethernet VPN (BGP EVPN). VXLAN is used as an overlay mechanism for carrying Layer 2 networks across a Layer 3 infrastructure, while BGP EVPN is used as the control plane for exchanging network information between Virtual Tunnel Endpoints (VTEPs). The implementation of these technologies provides more flexible network segmentation compared with conventional VLAN-based segmentation [7].

The network segmentation in this research consists of several VLANs mapped to specific VNIs. VLAN 3 is used for the DATA network and mapped to VNI 100003. VLAN 4 is used for the CCTV network and mapped to VNI 100004. VLAN 5 is used for the FIDS network and mapped to VNI 100005. VLAN 142 is used for the PERKANTORAN network and mapped to VNI 100142. In addition, VNI 110001 is used as the Layer 3 VNI for VRF INTERNAL. This mapping is used to maintain service segmentation within the overlay network [8].

The network architecture implemented in this research consists of two Spine devices, four Leaf devices, two Switch Access or STDT devices, one FortiGate device as the perimeter device, one Ubuntu Server, and several client devices. Two Leaf devices are configured as an MLAG pair. The FortiGate device functions as the perimeter device connecting the internal network to the external network. The Ubuntu Server functions as the service server, while the client devices are used as traffic sources during connectivity testing.

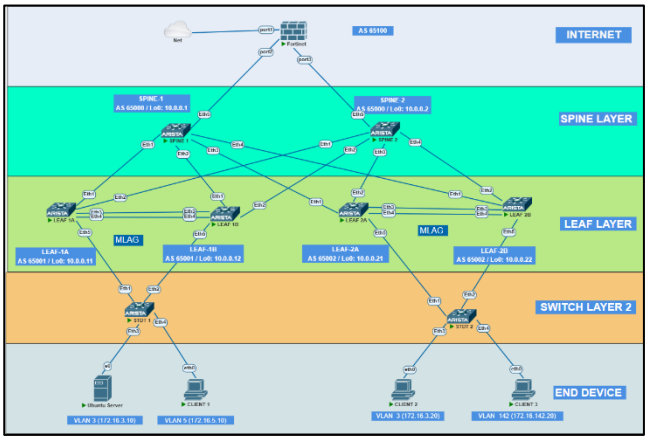

Figure 1 - Network topology implementation of the Terminal 3 Data Center

Figure 1 illustrates the network topology implemented in this research. The topology shows the relationships among the

Spine devices, Leaf devices, Switch Access devices, FortiGate, server, and client devices. The topology serves as the foundation for implementing the Leaf-Spine architecture, MLAG, VXLAN, BGP EVPN, Anycast Gateway, and ECMP.

The network implementation is conducted in a PNETLab simulation environment using Arista vEOS 4.31.0F virtual devices. The simulation environment is used because failure testing directly on the production network of Terminal 3 cannot be performed without potentially affecting operational services. The simulation environment allows the proposed network architecture to be implemented and tested without disrupting the existing production infrastructure.

In addition to the Arista devices, the testing environment uses FortiGate VM64-KVM as the perimeter device, Ubuntu Server as the application server, Uptime Kuma as the monitoring system, and Wireshark as the packet analysis tool. These components are used to perform network implementation, connectivity testing, failover testing, service availability monitoring, and verification of VXLAN packet encapsulation [9]. The Failover mechanism is evaluated through three failure scenarios. The first scenario represents an access-link failure by disabling Ethernet2 on T3-STDT-1. This scenario is used to evaluate the ability of MLAG to maintain connectivity when one of the access paths fails. The second scenario represents a complete Leaf device failure. In this scenario, Ethernet1, Ethernet2, and Ethernet5 on T3-LEAF-1A are disabled. This test is conducted to evaluate the ability of the Leaf pair to maintain connectivity when one Leaf device is no longer able to forward traffic. The third scenario represents a Spine device failure. In this scenario, Ethernet1 through Ethernet5 on T3-SPINE-1 are disabled. This test is conducted to evaluate the ability of ECMP to redirect traffic through the remaining active Spine device.

The evaluation is performed by sending ICMP traffic to three test destinations. The first destination is the Anycast Gateway of VLAN 3 with the address 172.16.3.254. The second destination is a client device on another network segment with the address 172.16.3.20. The third destination is the Internet at 8.8.8.8. These three destinations are used to evaluate connectivity at different network layers, ranging from the local gateway and the network fabric to external network connectivity [10].

Each test sends 150 ICMP packets to each destination, resulting in a total of 450 packets evaluated for each failure scenario. The observed parameters include packet loss, network recovery time, and service response time. Uptime Kuma is used to monitor FortiGate WAN connectivity, Internet connectivity to 8.8.8.8, and the T3 Web Server service. This testing approach is used to identify differences between network connectivity recovery and application service recovery.

VXLAN implementation is verified through packet capture using Wireshark. The verification is performed to confirm that traffic traversing the network fabric undergoes VXLAN encapsulation according to the proposed design. The analysis includes the VNI information, source and destination addresses in the outer header, and the VXLAN header structure observed in the captured packets.

The objective of this research is to design and implement a Data Center network based on the Leaf-Spine architecture using Arista devices with the implementation of VXLAN, BGP EVPN, MLAG, Anycast Gateway, and ECMP. In addition, this research aims to evaluate the ability of the Failover mechanism to maintain network accessibility when failures occur at the access-link, Leaf-device, and Spine-device levels.

The contribution of this research focuses on implementing redundancy mechanisms across multiple network layers. At the access layer, MLAG is implemented to provide Active-Active paths. At the Leaf layer, a redundant device pair with an Anycast Gateway is implemented to provide device-level redundancy. At the Spine layer, ECMP is implemented to provide alternative paths through the remaining Spine device. Meanwhile, VXLAN and BGP EVPN are implemented to support network segmentation and communication within the overlay network.The test results show that the Failover mechanism is able to maintain network connectivity across all three failure scenarios. All scenarios achieved 0% packet loss from 450 transmitted packets per scenario, with network recovery occurring in less than 200 milliseconds. During the Leaf device failure, the HTTP service experienced approximately 20 seconds of disruption, while the Spine device failure resulted in approximately 130 seconds of HTTP service disruption.The Spine device failure test revealed a difference between network recovery time and HTTP service recovery time. The Arista fabric detected the failure through the Hold Timer Expired mechanism after 9 seconds. Meanwhile, the difference in BGP Keepalive and Hold Time configurations between the Arista devices and the FortiGate device caused the service recovery process at the perimeter layer to take longer. This condition demonstrates that successful Failover convergence within the network fabric does not necessarily result in the same recovery time at the application-service layer [11].

Based on these results, this research evaluates not only the efectiveness of link and device redundancy but also identifies configuration factors that influence overall service recovery time. Through the implementation of the Leaf-Spine architecture, MLAG, VXLAN, BGP EVPN, Anycast Gateway, and ECMP, the proposed design is expected to provide an approach for improving accessibility and High Availability in Data Center networks that require continuous service availability.

2. Method

2.1 Research Method

This research uses an experimental simulation-based approach to design, implement, and evaluate a Data Center network with a Failover mechanism. The experiment is conducted in a controlled network simulation environment to evaluate the behavior of the proposed architecture under several failure conditions. The tested system consists of an Arista-based Leaf-Spine architecture supported by VXLAN, BGP EVPN, MLAG, Anycast Gateway, and ECMP. The simulation-based approach was selected because failure testing cannot be directly performed on the production network of Terminal 3 Soekarno-Hatta Airport due to its continuous support for airport operational services [12]. The simulation environment allows network failures to be introduced in a controlled manner without affecting the operational infrastructure. The research focuses on evaluating network recovery, packet loss, and service availability when link and device failures occur [13].

The research process consists of four main stages: requirements analysis, system design, implementation, and system testing. During the requirements analysis stage, the existing network conditions and redundancy requirements are identified. The system design stage defines the Leaf-Spine topology, network segmentation, routing architecture, and Failover mechanisms. The implementation stage builds the proposed network in PNETLab using virtual Arista devices. Finally, the testing stage evaluates the network under predefined failure scenarios.

2.2 Network Design and Implementation

The proposed network follows a two-layer Leaf-Spine architecture consisting of two Spine devices and four-Leaf devices. Each Leaf device is connected to both Spine devices, providing redundant paths within the network fabric. The Leaf devices are organized into two MLAG pairs to provide redundancy at the access layer.

The first MLAG pair consists of T3-LEAF-1A and T3-LEAF-1B, while the second pair consist of T3-LEAF-2A and T3-LEAF-2B. The access switches T3-STDT-1 and T3-STDT-2 are connected to their respective Leaf pairs through Port-Channel connections. This configuration allows both physical links to operate actively while maintaining connectivity when one link or Leaf device fails [14].

VXLAN is implemented as the overlay data plane, while BGP EVPN is used as the control plane for exchanging network reachability information between VTEPs. The network segments are mapped to VXLAN Network Identifiers (VNIs), consisting of VLAN 3 to VNI 100003 for DATA, VLAN 4 to VNI 100004 for CCTV, VLAN 5 to VNI 100005 for FIDS, and VLAN 142 to VNI 100142 for PERKANTORAN. VNI 110001 is used as the Layer 3 VNI for the INTERNAL VRF.

The routing fabric uses ECMP to provide multiple equal-cost paths between Leaf and Spine devices. Each Leaf is connected to both Spine devices through point-to-point Layer 3 links. The Spine devices also provide redundant connections toward the FortiGate perimeter device. This design allows traffic to use an alternative path when one of the Spine devices becomes unavailable.

The Anycast Gateway is configured on the Leaf pair to provide a common gateway function for connected network segments. The implementation also uses an Anycast VTEP address on the MLAG Leaf pair to maintain overlay connectivity when one member of the pair fails.

2.3 Experimental Environment

The proposed network is implemented in PNETLab using virtual Arista vEOS devices. FortiGate VM64-KVM is used as the perimeter firewall and border router, while Ubuntu Server provides the Web Server service used to evaluate service availability. Uptime Kuma is used to monitor network and HTTP availability, and Wireshark is used to verify VXLAN packet encapsulation.

The main components of the experimental environment are presented in Table 1.

Table 1 - Experimental environment specifications

No.	Component	Specification
1	Simulation platform	PNETLab
2	Network operating system	Arista vEOS 4.31.0F
3	Perimeter device	FortiGate VM64-KVM, FortiOS v6.0.2
4	Application server	Ubuntu Server, Python3 ThreadingTCPServer, port 8080
5	Monitoring system	Uptime Kuma
6	Test client	Linux host on VLAN 3 (DATA)
7	Packet analysis	Wireshark

The network environment contains two Spine devices, four Leaf devices, two access switches, one FortiGate perimeter device, one Ubuntu Server, and client devices. FortiGate is connected to both Spine devices through separate Layer 3 interfaces and establishes BGP sessions with the Spine layer. This configuration provides redundant connectivity between the internal Data Center fabric and the external network.

The network addressing uses private Autonomous System (AS) numbers. The Spine layer uses AS 65000, MLAG pair 1 uses AS 65001, MLAG pair 2 uses AS 65002, and the FortiGate perimeter device uses AS 65100. The Spine devices use loopback addresses 10.0.0.1 and 10.0.0.2, while the Leaf devices use 10.0.0.11 and 10.0.0.12 for the first MLAG pair. The first Leaf pair shares 10.0.1.1 as its Anycast VTEP address.

2.4 Failover Testing Procedure

The Failover evaluation is performed by intentionally disabling selected network interfaces and observing the resulting changes in connectivity and service availability. Three failure scenarios are defined to evaluate redundancy mechanisms at different network layers. The first scenario evaluates access-link redundancy by disabling Ethernet2 on T3-STDT-1. The second scenario evaluates Leaf-level redundancy by disabling Ethernet1, Ethernet2, and Ethernet5 on T3-LEAF-1A, effectively isolating the device from the network. The third scenario evaluates Spine-level redundancy by disabling Ethernet1 through Ethernet5 on T3-SPINE-1.

The three scenarios and the corresponding redundancy mechanisms are presented in Table 2.

Table 2 - Failover testing scenarios

Code	Failure Type	Failure Action	Mechanism Evaluated
S1	Access-link failure	Disable Ethernet2 on T3-STDT-1	MLAG at the access layer
S2	Leaf device failure	Disable Ethernet1, Ethernet2, and Ethernet5 on T3-LEAF-1A	Anycast VTEP at the Leaf layer
S3	Spine device failure	Disable Ethernet1–Ethernet5 on T3-SPINE-1	ECMP at the Spine layer

Before conducting the first scenario, traffic counters are checked to identify an active Uplink path. Ethernet2 on T3-STDT-1 is then selected as the failure target. This procedure ensures that

the link selected for failure is actively carrying traffic and that the test evaluates the actual Failover behavior of the MLAG connection.

For each scenario, the Failover test is performed using the Failover_test.sh script from a client located on VLAN 3. The script sends 150 ICMP packets to each test destination at an interval of 0.2 seconds. Three destinations are used to evaluate connectivity at different network layers: the VLAN 3 Anycast Gateway at 172.16.3.254, a client on another network segment at 172.16.3.20, and the Internet at 8.8.8.8. Therefore, a total of 450 ICMP packets are transmitted in each failure scenario.

The measured parameters consist of packet loss, network recovery time, round-trip time (RTT), and service availability. Packet loss and RTT are obtained from the ICMP test, while service availability is observed using Uptime Kuma. The monitoring system checks the FortiGate WAN interface at 192.168.149.144, Internet connectivity to 8.8.8.8, and the HTTP service provided by the T3 Web Server.

The Uptime Kuma ping monitors operate at a 20-second interval, while the HTTP monitor uses a 4000 ms timeout. The HTTP service is exposed through the FortiGate Virtual IP and forwarded to the Ubuntu Server on port 8080. This configuration allows the experiment to evaluate not only network-level connectivity but also the availability of an application service during Failover.

2.5 Data Analysis

The collected data are analyzed quantitatively by comparing network behavior before, during, and after each failure scenario. Packet loss is determined from the number of ICMP packets that fail to reach the destination, while network recovery time represents the observed interval required for connectivity to return after a failure is introduced.

RTT values are analyzed to identify temporary latency increases during the Failover process. Service availability data from Uptime Kuma are compared with the ICMP results to determine whether network convergence and application-service recovery occur at the same time. In addition to connectivity testing, Wireshark packet capture is used to verify the VXLAN implementation. The captured packets are examined to identify the VNI, outer source and destination addresses, and VXLAN encapsulation structure. This verification is used to confirm that the implemented overlay network operates according to the proposed design. The overall effectiveness of the proposed architecture is determined by comparing the results of the three failure scenarios. The analysis focuses on the ability of MLAG, Anycast VTEP, and ECMP to maintain connectivity, minimize packet loss, and support service continuity when failures occur at different network layers.

3. Result and Discussion

3.1 Redundant Path and Network Fabric Verification

The implemented network was first verified to ensure that the redundant paths between the Leaf and Spine layers were operating as designed. The underlay network uses point-to-point /31 links between each Leaf and both Spine devices. Therefore, each Leaf device has two independent paths toward the Spine layer, preventing a single physical link from becoming the only available path.

The verification results showed that the interfaces connecting T3-LEAF-1A to T3-SPINE-1 and T3-SPINE-2 were in an up/up state with an MTU of 9214 bytes. The MTU configuration was selected to accommodate the additional overhead introduced by VXLAN encapsulation. Each Leaf device also established separate BGP sessions with both Spine devices, confirming the availability of redundant paths in the control plane.

At the access layer, the MLAG configuration allowed the two physical Uplink connections from the access switch to operate as a single logical Port-Channel. When both links were active, traffic could be forwarded through both paths. Unlike the previous network condition, in which part of the Uplink capacity was blocked by the Spanning Tree mechanism, the implemented design allowed both Uplinks to remain in the forwarding state.

The routing verification also showed two active next-hop paths through the Spine layer. The redundant paths were represented by VTEP addresses 10.0.2.1 and 10.0.2.2. This configuration confirms that the network fabric can provide an alternative path when one Spine device becomes unavailable.

These results demonstrate that the proposed Leaf-Spine architecture provides path redundancy at both the access and fabric layers. MLAG provides redundancy between the access switches and Leaf pair, while ECMP provides multiple equal-cost paths between the Leaf and Spine layers.

3.2 VXLAN and BGP EVPN Verification

The implementation of VXLAN and BGP EVPN was verified through both device-level configuration and packet capture. The configured service segments consisted of DATA, CCTV, FIDS, and PERKANTORAN, which were mapped to their respective VXLAN Network Identifiers. VNI 110001 was used as the Layer 3 VNI for the INTERNAL VRF.

The BGP EVPN verification showed that MAC-IP information was distributed between VTEPs through the EVPN control plane. The presence of EVPN-learned MAC addresses indicates that host information was distributed through the control plane rather than relying only on conventional Layer 2 flooding.

Packet capture was subsequently performed on the link between T3-LEAF-1A and T3-SPINE-1 to verify that the configured VXLAN overlay was actually applied to network traffic. The captured packet consisted of an outer Ethernet header, an outer IPv4 header, UDP, VXLAN, and the original inner Ethernet and IPv4 information.

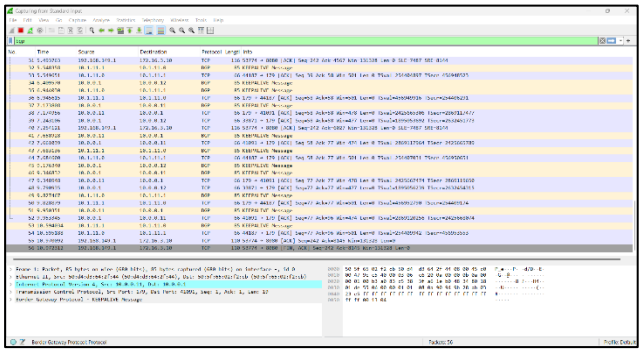

Figure 2 - VXLAN encapsulation decomposition from packet capture

The packet analysis identified an outer Ethernet header of 14 bytes, an outer IPv4 header of 20 bytes, an 8-byte UDP header, and an 8-byte VXLAN header. Therefore, the additional encapsulation overhead was 50 bytes. The VXLAN header contained VNI 110001, confirming that the captured traffic belonged to the Layer VNI associated with the INTERNAL VRF. The outer source address was 10.0.1.1, representing the Anycast VTEP of MLAG-PAIR1, while the outer destination address was 10.0.2.1, representing the Border VTEP on T3-SPINE-1. The result confirms that the traffic was transported through the VXLAN overlay between the Anycast VTEP and the Border VTEP.

The 50-byte overhead also supports the use of an MTU of 9214 bytes on the underlay links. The corresponding system VLAN used for the Layer 3 VNI showed an MTU of 9164 bytes, which is consistent with the 50-byte VXLAN overhead.
Thus, the VXLAN and BGP EVPN implementation was verified not only from the configuration output but also at the packet level.

3.3 Failover Test Results

Three failure scenarios were conducted to evaluate the ability of the proposed architecture to maintain connectivity during network failures. Each scenario transmitted 150 ICMP packets to three destinations: the VLAN 3 Anycast Gateway at 172.16.3.254, a client at 172.16.3.20, and the Internet at 8.8.8.8. Consequently, 450 packets were evaluated in each scenario.

a. Access-Link Failure

The first scenario evaluated MLAG redundancy by disabling Ethernet2 on T3-STDT-1. After the interface was disabled, Ethernet1 remained active and the Port-Channel continued operating with one active member.

The test results showed that all 150 packets were successfully delivered to each destination. No packet loss occurred for the gateway, client, or Internet destination. The average RTT values were 7.275 ms for the gateway, 30.993 ms for the client, and 67.329 ms for the Internet. The maximum RTT reached 97.096 ms, 403.459 ms, and 236.708 ms, respectively. The network recovery time was measured at less than 200 milliseconds. The temporary increase in RTT, particularly the maximum RTT of 403.459 ms toward the client, indicates the transient effect of traffic path switching during Failover. However, the absence of packet loss demonstrates that MLAG successfully maintained connectivity through the remaining active Uplink.

b. Leaf Device Failure

The second scenario evaluated Leaf-level redundancy by completely isolating T3-LEAF-1A through the disabling of Ethernet1, Ethernet2, and Ethernet5. Ethernet1 and Ethernet2 were the connections toward the Spine layer, while Ethernet5 was associated with the access connection.

Despite the complete failure of T3-LEAF-1A, all 150 packets sent to each destination were successfully received. The average RTT values were 12.590 ms for the gateway, 24.843 ms for the client, and 57.489 ms for the Internet. The maximum RTT values were 356.717 ms, 285.745 ms, and 316.610 ms, respectively. The network recovery time remained below 200 milliseconds, with 0% packet loss. The increase in gateway RTT compared with the first scenario reflects the additional traffic load handled by the surviving Leaf device.

Uptime Kuma showed that the FortiGate WAN and Internet monitors remained available at 100%. However, the T3 Web Server experienced a short service disruption of approximately 20 seconds during the transition from T3-LEAF-1A to T3-LEAF-1B. This indicates that the network fabric-maintained connectivity while the application-level service experienced a temporary disruption during the MLAG switchover.

c. Spine Device Failure

The third scenario evaluated ECMP redundancy by disabling Ethernet1 through Ethernet5 on T3-SPINE-1. This action isolated T3-SPINE-1 from the network fabric and forced the remaining traffic to use T3-SPINE-2. The test results again showed 0% packet loss for all three destinations. The average RTT values were 8.096 ms for the gateway, 21.513 ms for the client, and 67.535 ms for the Internet. The maximum RTT values were 172.612 ms, 215.190 ms, and 301.926 ms, respectively.

The network recovery time remained below 200 milliseconds. The results demonstrate that the ECMP-based fabric was able to redirect traffic through the remaining Spine device without packet loss. However, the application-level result differed from the ICMP connectivity result. Uptime Kuma recorded that the T3 Web Server became unavailable at 01:36:24 with a 4000 ms timeout and returned to normal operation at 01:38:34 with an HTTP 200 OK response. The observed HTTP service disruption was therefore 130 seconds.

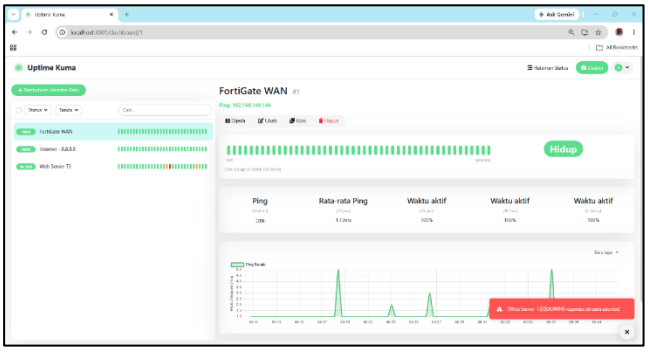

Figure 3 - Service event record during the Spine failure scenario

The monitoring results showed that the FortiGate WAN and Internet monitors remained at 100% availability, while only the T3 Web Server monitor experienced an interruption (Figure 3). This indicates that the perimeter device and outbound Internet path remained operational, while the problem occurred on the inbound path toward the internal Web Server.

3.4 Summary of Failover Results

The results of the three failure scenarios are summarized in Table 3.

Table 3 - Summary of Failover test results

Scenario	Failure Type	Packet Loss	Network Recovery	HTTP Service Disruption
S1	Access-link failure	0%	<200 ms	Not recorded
S2	Leaf device failure	0%	<200 ms	~20 s
S3	Spine device failure	0%	<200 ms	130 s

The results show that all three failure scenarios-maintained network connectivity with 0% packet loss from 450 transmitted packets per scenario. Network recovery was consistently measured at less than 200 milliseconds. This demonstrates that the redundancy mechanisms implemented at the access, Leaf, and Spine layers were able to maintain network-level connectivity during the tested failures.

The differences observed at the HTTP service level demonstrate that network recovery and application-service recovery do not necessarily occur at the same time. The access-link failure did not produce a measurable HTTP interruption, while the Leaf failure produced an approximately 20-second disruption. The longest disruption occurred during the Spine failure, although the ICMP tests showed that the network fabric itself recovered in less than 200 milliseconds.

3.5 Analysis of the Spine Failure and BGP Timer

The 130-second HTTP disruption observed during the Spine failure required further analysis because the ICMP test showed 0% packet loss and a network recovery time below 200 milliseconds. The analysis therefore focused on the BGP relationship between the Arista fabric and the FortiGate perimeter device.

The BGP configuration showed a significant difference between the timers used by the two platforms. Arista vEOS used a Keepalive interval of 3 seconds and a Hold Time of 9 seconds, while FortiGate used a Keepalive interval of 60 seconds and a Hold Time of 180 seconds.

Table 4 - BGP timer configuration comparison

Device	Keepalive	Hold Time	Configuration Source
Arista vEOS (fabric)	3 s	9 s	Explicit configuration
FortiGate (perimeter)	60 s	180 s	FortiOS default value

The Arista system log recorded a Hold Timer Expired event 9 seconds after T3-SPINE-1 was disabled. This confirms that the Arista fabric detected the failure and converged through the remaining Spine path significantly faster than the application service recovered. The difference became more apparent when the monitoring events were reconstructed. The Spine interface was disabled at 01:36:08. The Arista device declared the BGP session down at 01:36:17, corresponding to a 9-second detection interval. The HTTP service was subsequently reported as unavailable at 01:36:24 and returned with an HTTP 200 OK response at 01:38:34. The analysis indicates that the prolonged HTTP disruption was caused by the BGP timer behavior on the FortiGate perimeter device rather than by the Arista fabric. The FortiGate continued to consider the affected BGP neighbor active until its Hold Time expired. Once the session was declared inactive, the route through T3-SPINE-1 was removed and traffic could use T3-SPINE-2.

This finding demonstrates that fast convergence within the Data Center fabric alone is not sufficient to guarantee fast application-service recovery. The recovery behavior also depends on the convergence mechanisms and timer configurations of devices located outside the fabric.

3.6 Overall Discussion

The experimental results demonstrate that the proposed architecture improves Data Center accessibility by introducing redundancy at multiple network layers. The previous single Backbone path was replaced with redundant paths through two Spine devices, while the access layer was provided with active redundant Uplinks through MLAG.At the fabric layer, ECMP successfully maintained connectivity when T3-SPINE-1 was disabled. At the Leaf layer, Anycast VTEP and MLAG allowed connectivity to continue when T3-LEAF-1A was completely isolated. At the access layer, MLAG maintained connectivity when one of the physical Uplink links was disabled.

The VXLAN and BGP EVPN implementation also successfully provided network segmentation across the fabric. The packet capture confirmed VNI 110001 and the 50-byte VXLAN encapsulation overhead, providing packet-level evidence that the overlay mechanism was operating according to the proposed design.

The most important finding is the distinction between network convergence and application-service recovery. Although all three scenarios achieved less than 200 milliseconds of network recovery and 0% packet loss, application-level disruptions were still observed in the Leaf and Spine failure scenarios. The Spine failure produced the longest disruption because of the different BGP timer configurations between the Arista fabric and FortiGate perimeter device. Therefore, the effectiveness of a High Availability architecture should not be evaluated solely from packet loss or network convergence. Monitoring should also include application-service availability and the behavior of adjacent network domains. In this implementation, aligning the BGP timers between the Arista fabric and FortiGate, or using a faster failure-detection mechanism such as Bidirectional Forwarding Detection (BFD), would reduce the delay between fabric convergence and service recovery.

4. Conclusion

The implementation of a redundant data center network using an Arista-based Spine-Leaf architecture, VXLAN, BGP EVPN, MLAG, Anycast VTEP, and ECMP successfully improved network accessibility and path redundancy. The proposed design eliminated the dependency on a single backbone path and enabled multiple active forwarding paths between the access layer and the Spine layer. The failover evaluation demonstrated that all three failure scenarios—access-link failure, Leaf device failure, and Spine device failure—achieved 0% packet loss, with network

recovery measured at less than 200 ms. These results indicate that the network fabric was able to maintain connectivity during link and device failures through redundant paths and dynamic route convergence. The Leaf failure scenario resulted in an approximately 20-second HTTP service disruption, while the Spine failure scenario resulted in a 130-second HTTP disruption. Further analysis showed that the prolonged disruption in the Spine failure scenario was not caused by the Arista fabric, which detected the failure within 9 seconds, but by the different BGP timer configurations between the Arista devices and the FortiGate perimeter device. Therefore, aligning the BGP timers or implementing BFD is recommended to reduce service recovery time and improve end-to-end failover performance. Overall, the proposed network design provides improved redundancy, utilization of multiple active links, and resilience against network failures, making it a suitable approach for improving data center accessibility and availability.

Acknowledgements

The author would like to express sincere gratitude to Universitas Pamulang, particularly the Faculty of Computer Science and the Informatics Engineering Study Program, for providing the academic environment and support throughout this research. The author also gratefully acknowledges the guidance and valuable assistance provided by the academic supervisor during the research, implementation, testing, and preparation of this article.
4. The author also thanks all parties who contributed to the implementation and evaluation of the proposed network architecture, particularly those who provided technical support during the simulation and testing process.

REFERENCES

[1] A. B. Pratomo, "PERBANDINGAN METODE FAILOVER RSTP DAN ROUTING POLICY PADA ROUTER MIKROTIK," *Bull. Netw. Eng. Informatics*, vol. 2, no. 1, 2024.

[2] Y. B. Ginting, U. Ristian, J. Rekayasa, and S. Komputer, "Implementasi Metode Failover Sebagai Backup Server," *Coding J. Komput. dan Apl.*, vol. 09, no. 02, 2021.

[3] Y. Pribadi, A. B. Putra Negara, and M. A. Irwansyah, "Analisis Penggunaan Metode Failover Clustering untuk Mencapai High Availability pada Web Server (Studi Kasus: Gedung Jurusan Informatika)," *J. Sist. dan Teknol. Inf.*, vol. 8, no. 2, 2020.

[4] M. N. Afdhol. P. Y., A. Anggraini Samudra, and R. Trisetyowati Untari, "PERANCANGAN JARINGAN KOMPUTER MENGGUNAKAN METODE FAILOVER," *JATI (Jurnal Mhs. Tek. Inform.*, vol. 7, no. 3, 2023.

[5] F. Guo and A. Ye, "The application and performance optimization of multi-controller-based load balancing algorithm in computer networks," *Egypt. Informatics J.*, vol. 30, 2025, doi: https://doi.org/10.1016/j.eij.2025.100678.

[6] A. B. Pratomo, "PENGEMBANGAN SISTEM FIREWALL PADA JARINGAN KOMPUTER BERBASIS MIKROTIK ROUTEROS DEVELOPING A FIREWALL SYSTEM ON A COMPUTER NETWORK BASED ON MIKROTIK ROUTEROS," *Bul. Netw. Informatics*, vol. 1, no. 2, 2023.

[7] Y. Peng, "Research on the Technology of Computer Network Security Protection," *J. Appl. Data Sci.*, vol. 4, no. 1, 2023, doi: https://doi.org/10.47738/jads.v4i1.80.

[8] A. Anton and A. Irman, "Implementasi Load Balance Mikrotik Dual ISP Dengan PCC dan Metode Failover Pada PT. Wahana Ciptasinatria," *J. Teknol. Inf.*, vol. 10, no. 1, 2024.

[9] N. Nirsal, R. Rusmala, F. E. Susilawati, S. Syafriadi, J. Y. Sari, and R. Sari, "Integrating Augmented Reality in Computer Network Practicum: A Comprehensive Review and Proposal for Enhanced Learning Experience," *TEM J.*, vol. 14, no. 2, 2025.

[10] Q. Liu and T. Zhang, "Deep learning technology of computer network security detection based on artificial intelligence," *Comput. Electr. Eng.*, vol. 110, 2023, doi: https://doi.org/10.1016/j.compeleceng.2023.108813.

[11] A. Novokhrestov, A. Konev, and A. Shelupanov, "Model of threats to computer network software," *Symmetry (Basel).*, vol. 11, no. 12, 2019.

[12] I. R. N. Siringo-Ringo, E. Veronika, M. Wahidin, and D. A. Kusumaningtiar, "THE VALUE OF AIR GERM NUMBERS AT SOEKARNO HATTA AIRPORT IN 2021," *J. Kesehat. Lingkung.*, vol. 14, no. 4, 2022.

[13] J. Wang and X. Wang, "An optimization model of computer network security based on GABP neural network algorithm," *Eurasip J. Inf. Secur.*, vol. 2025, no. 1, 2025, doi: https://doi.org/10.1186/s13635-025-00196-5.

[14] A. BaniMustafa, M. Baklizi, and K. Khatatneh, "Machine Learning for Securing Traffic in Computer Networks," *Int. J. Adv. Comput. Sci. Appl.*, vol. 13, no. 12, 2022, doi: https://doi.org/10.14569/IJACSA.2022.0131252.